



Department of Electrical Engineering
School of Science and Engineering

EE514/CS535 Machine Learning

HOMEWORK 1

Due Date: 23:55, Tuesday, March 08, 2022 (Submit online on LMS)

Format: 6 problems, for a total of 100 marks

Instructions:

- Each student must submit his/her own hand-written assignment, scanned in a **single PDF document**.
 - You are allowed to collaborate with your peers but copying other's solution is strictly prohibited. Anybody found guilty would be subjected to disciplinary action in accordance with the university rules and regulations.
 - Note: Vectors are written in lowercase and bold in the homework. For your written submission, kindly use an underline instead. In addition, use capital letters for matrices and lowercase for scalars.
-

Problem 1 (30 marks)

(Note: Compile and submit screenshots of your plots for this question.)

Curse of Dimensionality - In this question we will look closely at the problem associated with higher dimensions using k -NN. During your lectures, you learnt that as the dimensions increase, the neighbours in a constant vicinity decrease. Let us now visualise this phenomenon through simulations.

We are interested in finding the average distance between the two points randomly (uniformly) distributed inside a unit ball. d -dimensional unit ball of radius R is mathematically defined as

$$\mathbb{B}(d) = \{\mathbf{x} \in \mathbb{R}^d, \|\mathbf{x}\|_2 \leq R\},$$

which is referred to as unit ball for $R = 1$.

- (a) [2 marks] What does the the ball $\mathbb{B}(d)$ represent for $d = 2$ and $d = 3$?
- (b) [6 marks] You are required to build a function that will generate random (uniformly distributed) n number of data points inside the unit ball in d dimensional space. Generate a dataset where $n=50$ and $d=100$. Submit your code or attach a screenshot of your code.
- (c) [10 marks] You are required to build a function that will randomly take two points from a d -dimensional unit ball and calculate the Euclidean distance for $i = 10000$ iterations. You are given the following starter code:

```
import matplotlib.pyplot as plt
%matplotlib inline
import math
import random

def freq_plot(d_dim, iterations = 10000):
    dist = []
    # code required here

    plt.figure()
    plt.hist(dist, range = [0, math.sqrt(d_dim)], bins=100,
             density = True)
    title = 'n = ' + str(d_dim)
    plt.gca().set(title=title, xlabel='Distances', ylabel='Frequency')
    return
```

Use the function developed in part (b) for randomly selecting two points inside the unit ball.

Attach a screenshot of your final working function code.

- (d) [6 marks] Run the function in part (c) for $d_dim = 2, 3, 5, 10, 100$, and attach screenshots of your plots.
- (e) [6 marks] Interpret your results by explaining the difference in the plots for each dimension.

Problem 2 (20 marks)

Principal Component Analysis - In class we studied PCA as a means to reduce dimensionality, by minimizing the squared error of predictions in lower dimension, also known as projections onto a range space, as compared to the actual values. In this question, we will prove that this problem is equivalent to maximizing the squared length of these projections, i.e., maximizing variance of projected data.

Let us define the original problem from the class lecture. Considering n feature vectors of the form $\mathbf{x} \in \mathbb{R}^d$. Using only k basis vectors, we want to project \mathbf{x} in a new space of lower dimensionality, from \mathbb{R}^d to \mathbb{R}^k :

$$\mathbf{z} = \mathbf{W}^T \mathbf{x},$$

where \mathbf{W} is the orthogonal mapping matrix of size $d \times k$. To obtain an approximation of \mathbf{x} we use the following reconstruction:

$$\hat{\mathbf{x}} = \mathbf{W} \mathbf{z},$$

which is equivalent to:

$$\hat{\mathbf{x}} = \mathbf{W} \mathbf{W}^T \mathbf{x}$$

$$\hat{\mathbf{x}} = \mathbf{P} \mathbf{x},$$

where $\mathbf{P} = \mathbf{W} \mathbf{W}^T$ is the projection operator, that is idempotent: has the following properties: $\mathbf{P}^2 = \mathbf{P} = \mathbf{P}^T$.

Our objective is to minimize the sum of squared error during reconstruction, we can write it as follows:

$$\text{minimize } \|\mathbf{x} - \mathbf{P} \mathbf{x}\|_2^2$$

(a) [10 marks] Show that this expression is minimized when:

$$\mathbf{P} = \mathbf{W} \mathbf{W}^T$$

(b) [10 marks] Perform PCA on the following input matrix \mathbf{X} to reduce the number of dimensions to $k = 1$, where $n = 6$ and $d = 2$, such that $\mathbf{x} \in \mathbb{R}^d$ and $\hat{\mathbf{x}} \in \mathbb{R}^k$.

$$\mathbf{X} = \begin{bmatrix} 2 & 4 & 5 & 5 & 3 & 2 \\ 2 & 3 & 4 & 5 & 4 & 3 \end{bmatrix}$$

Problem 3 (10 marks)

Analytical Solution of Least Square - Linear regression can be expressed in Matrix notation:

$$\mathbf{y} = \mathbf{X} \boldsymbol{\theta}$$

Here $\mathbf{X} \in \mathbb{R}^{n \times m}$ (where $n > m$) is the input data and each column is a data feature, $\boldsymbol{\theta} \in \mathbb{R}^m$ is a vector of coefficients, and $\mathbf{y} \in \mathbb{R}^n$ is a vector of output variables for each row in \mathbf{X} .

$$\begin{bmatrix} y_1 \\ y_2 \\ \vdots \\ y_N \end{bmatrix} = \begin{bmatrix} 1 & x_1 & x_1^2 & \dots & x_1^M \\ 1 & x_2 & x_2^2 & \dots & x_2^M \\ \vdots & \vdots & \vdots & \ddots & \vdots \\ 1 & x_N & x_N^2 & \dots & x_N^M \end{bmatrix} \begin{bmatrix} \theta_0 \\ \theta_1 \\ \vdots \\ \theta_M \end{bmatrix}$$

This is an over-determined system which results in multiple possible values for the coefficients. A typical way to find a solution is where the values of $\boldsymbol{\theta}$ in the model minimize the squared error (linear least square solution). Recall that the least square expression is $\|\mathbf{X}\boldsymbol{\theta} - \mathbf{y}\|_2^2$. In matrix notation, this problem is formulated using the so-called normal equation that can be rearranged to give:

$$\boldsymbol{\theta} = (\mathbf{X}^T \mathbf{X})^{-1} \mathbf{X}^T \mathbf{y}$$

- (a) Given the following \mathbf{X} and \mathbf{y} , find the least square solution for $\boldsymbol{\theta}$.
(**Note:** Submit the code and final result for this part)

$$\mathbf{X} = \begin{bmatrix} 0.05 \\ 0.18 \\ 0.31 \\ 0.42 \\ 0.5 \end{bmatrix}, \quad \mathbf{y} = \begin{bmatrix} 0.12 \\ 0.22 \\ 0.35 \\ 0.38 \\ 0.49 \end{bmatrix}$$

- (b) Draw a scatter plot of \mathbf{y} vs \mathbf{X} . On the same plot, draw a line plot for the model using the value of $\boldsymbol{\theta}$ found in the previous part.
(**Note:** Compile and submit screenshots of your plots for this question.)

Problem 4 (10 marks)

Pakistan International Airlines has developed 2 different classifiers (A and B) for the prediction whether a flight originating from Lahore will arrive at its final destination on time or not. True or Positive here is 'On time' and it refers to the case when the flight is no more than 5 minutes late than the scheduled time. The classifiers were tested on a data-set of 500 flights, and the results are as follows:

	Actual	
	On time	Late
Classifier A, predicted on time	131	155
Classifier A, predicted late	19	195
Classifier B, predicted on time	82	72
Classifier B, predicted late	68	278

- (a) [5 marks] Which is the preferable classifier in terms of F_1 score?
- (b) [5 marks] Which is the preferable classifier in terms of accuracy?

Problem 5 (15 marks)

Multivariate Gradient Descent - Gradient descent is an iterative optimization algorithm used to find a local minimum of a function. You are presented with the following feature vector \mathbf{X} , the parameter vector $\boldsymbol{\theta}$, the bias term \mathbf{b} and the gold label \mathbf{y} .

$$\mathbf{X} = \begin{bmatrix} x_1 \\ x_2 \\ x_3 \end{bmatrix} = \begin{bmatrix} 4 \\ 9 \\ 5 \end{bmatrix}, \quad \boldsymbol{\theta} = \begin{bmatrix} \theta_1 \\ \theta_2 \\ \theta_3 \end{bmatrix} = \begin{bmatrix} -0.8 \\ 1.0 \\ 0.3 \end{bmatrix}, \quad \mathbf{b} = 0.5, \quad \mathbf{y} = 2$$

- (a) [13 marks] Perform 3 iterations of batch gradient descent on the dataset to find updated parameter vectors. Assume learning rate $\alpha = 0.001$ and round your answers to 4 decimal places. Use the least square loss function, and show all your steps. Verify your answer by showing that the loss function is decreasing.
- (b) [2 marks] Why is the learning rate usually a small value? What is the caveat of a really small learning rate?

Problem 6 (15 marks)

Weighted Ridge Regression - Suppose we have a dataset $\{(\mathbf{x}_1, y_1), \dots, (\mathbf{x}_n, y_n)\}$. Unlike ridge regression where each example is equally important, sometimes certain examples are more important than others and we would want to assign a positive weight ω_i to each training example to indicate the level of importance of each training example. The corresponding loss function is defined as

$$L(\boldsymbol{\beta}) = \sum_{i=1}^N \omega_i (\boldsymbol{\beta}^T \mathbf{x}_i - y_i)^2$$

- (a) [5 marks] Suppose $\mathbf{X} = [\mathbf{x}_1, \dots, \mathbf{x}_n]^T$ and $\mathbf{Y} = [y_1, \dots, y_n]^T$. Find a diagonal matrix \mathbf{W} such that we can rewrite the loss function as

$$L(\boldsymbol{\beta}) = (\mathbf{Y} - \mathbf{X}\boldsymbol{\beta})^T \mathbf{W} (\mathbf{Y} - \mathbf{X}\boldsymbol{\beta})$$

- (b) [10 marks] Using the rewritten loss function, derive a closed form solution for $\boldsymbol{\beta}$ by setting the gradient of the loss function equal to zero.

— End of Homework —